

IN THE CLAIMS

1. - 34. (Cancelled)

35. (New) An automated method for allocating resources among a plurality of resource-using computational entities in a data processing system, the method comprising:

establishing a service-level utility for each of said plurality of resource-using entities, wherein the service-level utility is representative of an amount of business value obtained by each of said plurality of resource-using entities for one or more levels of performance and demand associated with each resource-using entity;

transforming said service-level utility into a resource-level utility for each of said plurality of resource-using entities, wherein the resource-level utility is representative of an amount of business value obtained by each of said plurality of resource-using entities when a quantity of said resources is allocated to the resource-using entity, wherein the resource-level utility indicates, for at least one of said plurality of resource-using entities, an estimated cumulative discounted or undiscounted future utility starting from current state descriptions of said at least one resource-using entity, wherein said estimated cumulative discounted or undiscounted future utility is trained on a temporal sequence of observed data using an adaptive machine learning procedure;

aggregating said resource-level utilities of all of said plurality of resource-using entities;

computing a resource allocation from aggregated utility information by executing an optimization method to maximize a total utility of said data processing system, wherein said optimization method comprises a standard linear or nonlinear algorithm; and

executing and conveying to the plurality of resource-using entities said resource allocation.